INSTITUTE OF INFORMATICS
SLOVAK ACADEMY OF SCIENCES

A State-of-the-Art Local Training Methods in Federated Learning

Michal Staňo, Ladislav Hluchý

INSTITUTE OF INFORMATICS
SLOVAK ACADEMY OF SCIENCES, BRATISLAVA SLOVAKIA

IEEE 23rd International Symposium on Computational Intelligence and Informatics (CINTI 2023)
November 20, 2023

Outline of the Talk

1. What is Federated Learning?
2. What is Local Training
3. Brief History of Local Training
4. What does Local Training do?

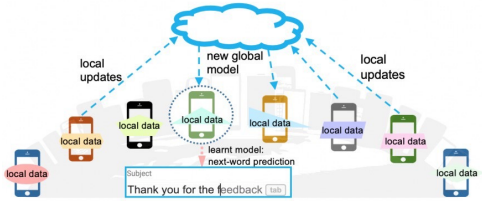

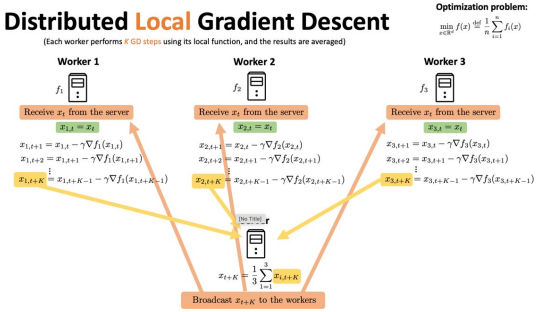

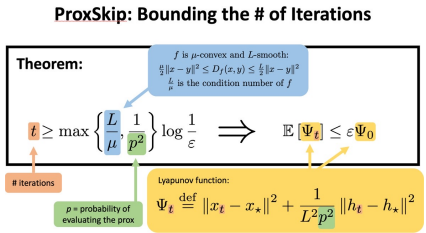

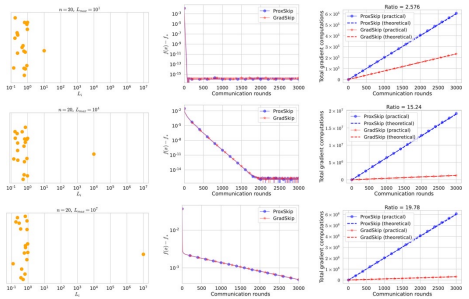

Algorithm 1 ProxSkip

stepsize $\gamma > 0$, probability $p > 0$, initial iterate $x_0 \in \mathbb{R}^d$, initial control variate $h_0 \in \mathbb{R}^d$, number of iterations $T \geq 1$
for $t = 0, 1, \ldots, T-1$ **do**
  $\hat{x}_{t+1} = x_t - \gamma(\nabla f(x_t) - h_t)$ ◇ Take a gradient-type step adjusted via the control variate h_t
  Flip a coin $\theta_t \in \{0, 1\}$ where $\text{Prob}(\theta_t = 1) = p$ ◇ Flip a coin that decides whether to skip the prox or not
  if $\theta_t = 1$ **then**
    $x_{t+1} = \text{prox}_{\frac{\gamma}{p}\psi}(\hat{x}_{t+1} - \frac{\gamma}{p} h_t)$ ◇ Apply prox, but only very rarely! (with small probability p)
  else
    $x_{t+1} = \hat{x}_{t+1}$ ◇ Skip the prox!
  end if
  $h_{t+1} = h_t + \frac{p}{\gamma}(x_{t+1} - \hat{x}_{t+1})$ ◇ Update the control variate h_t
end for

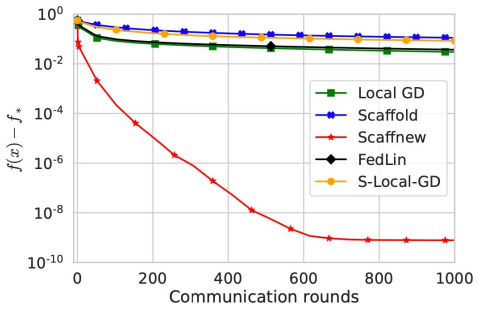

(c) theoretical hyper-parameters

Part 1
What is Federated Learning?

The First Federated Learning App: Next-Word Prediction?

Federated Learning is a collaborative machine learning from private data stored across a (large) number of clients/devices (e.g., hospitals, phones, banks)

Part 2
What is Local Training?

Local Training

A. Gradient Descent

B. Distributed Gradient Descent

C. Distributed Local Gradient Descent

Gradient Descent

$$\min_{x \in R^d} f(x)$$

$$x_{t+1} = x_t - \gamma \nabla f(x_t)$$

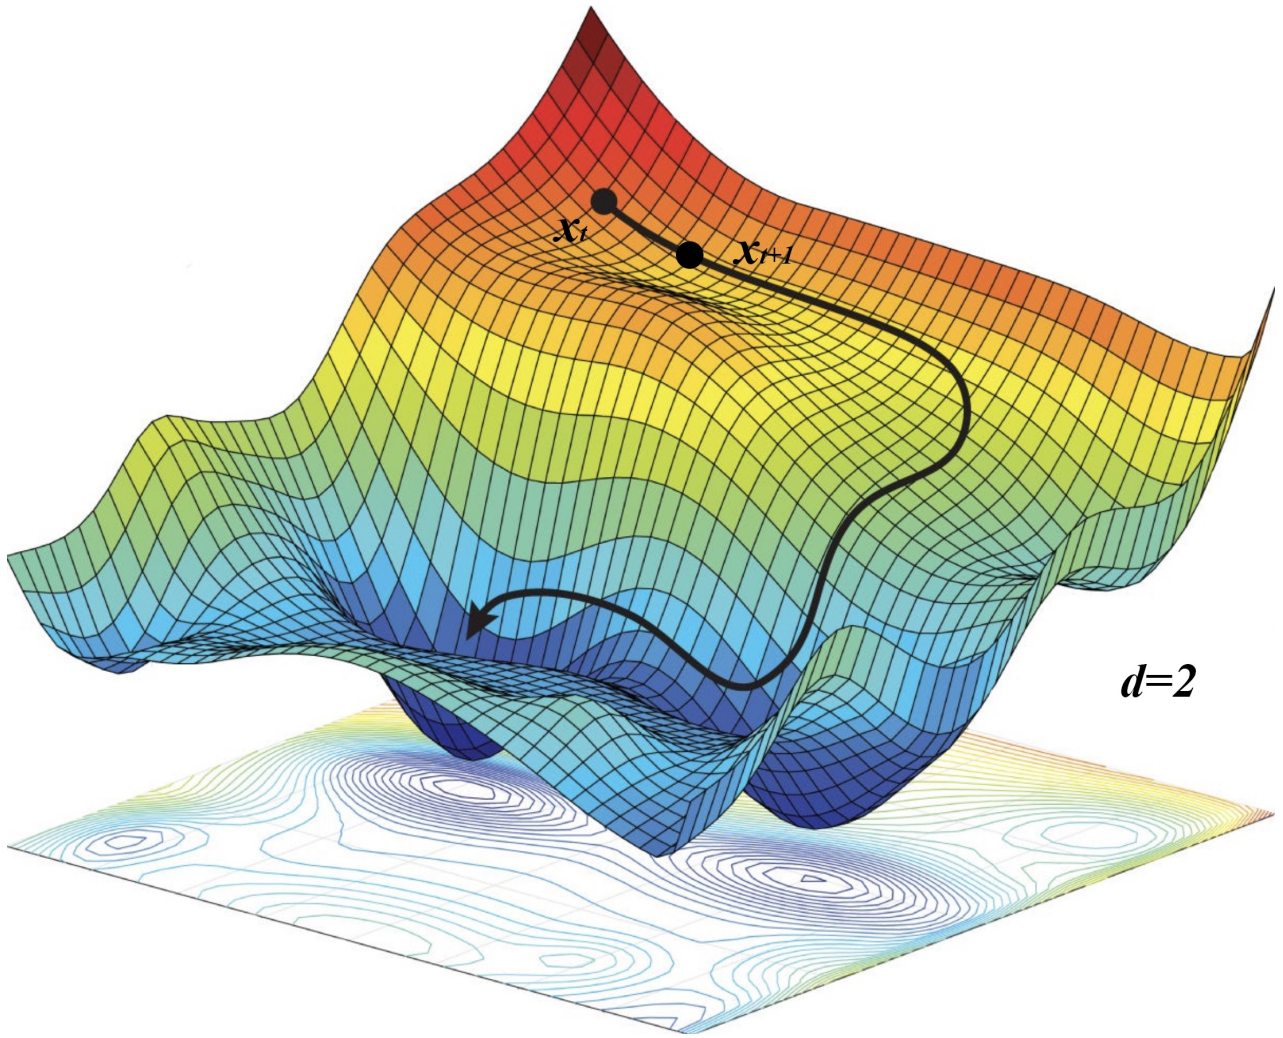

Distributed Gradient Descent

$$\min_{x \in R^d} f(x) \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^{n} f_i(x)$$

Distributed Gradient Descent

$$\min_{x \in \mathbb{R}^d} f(x) \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^{n} f_i(x)$$

n ← # devices / machines

d ← # model parameters / features

$f_i(x)$ ← Loss on local data D_i stored on device i

$$f_i(x) = \mathsf{E}_{\zeta \sim D_i} f_i(x, \zeta)$$

The datasets $D_1, \ldots, D_n$ can be arbitrarily heterogeneous

Distributed Gradient Descent

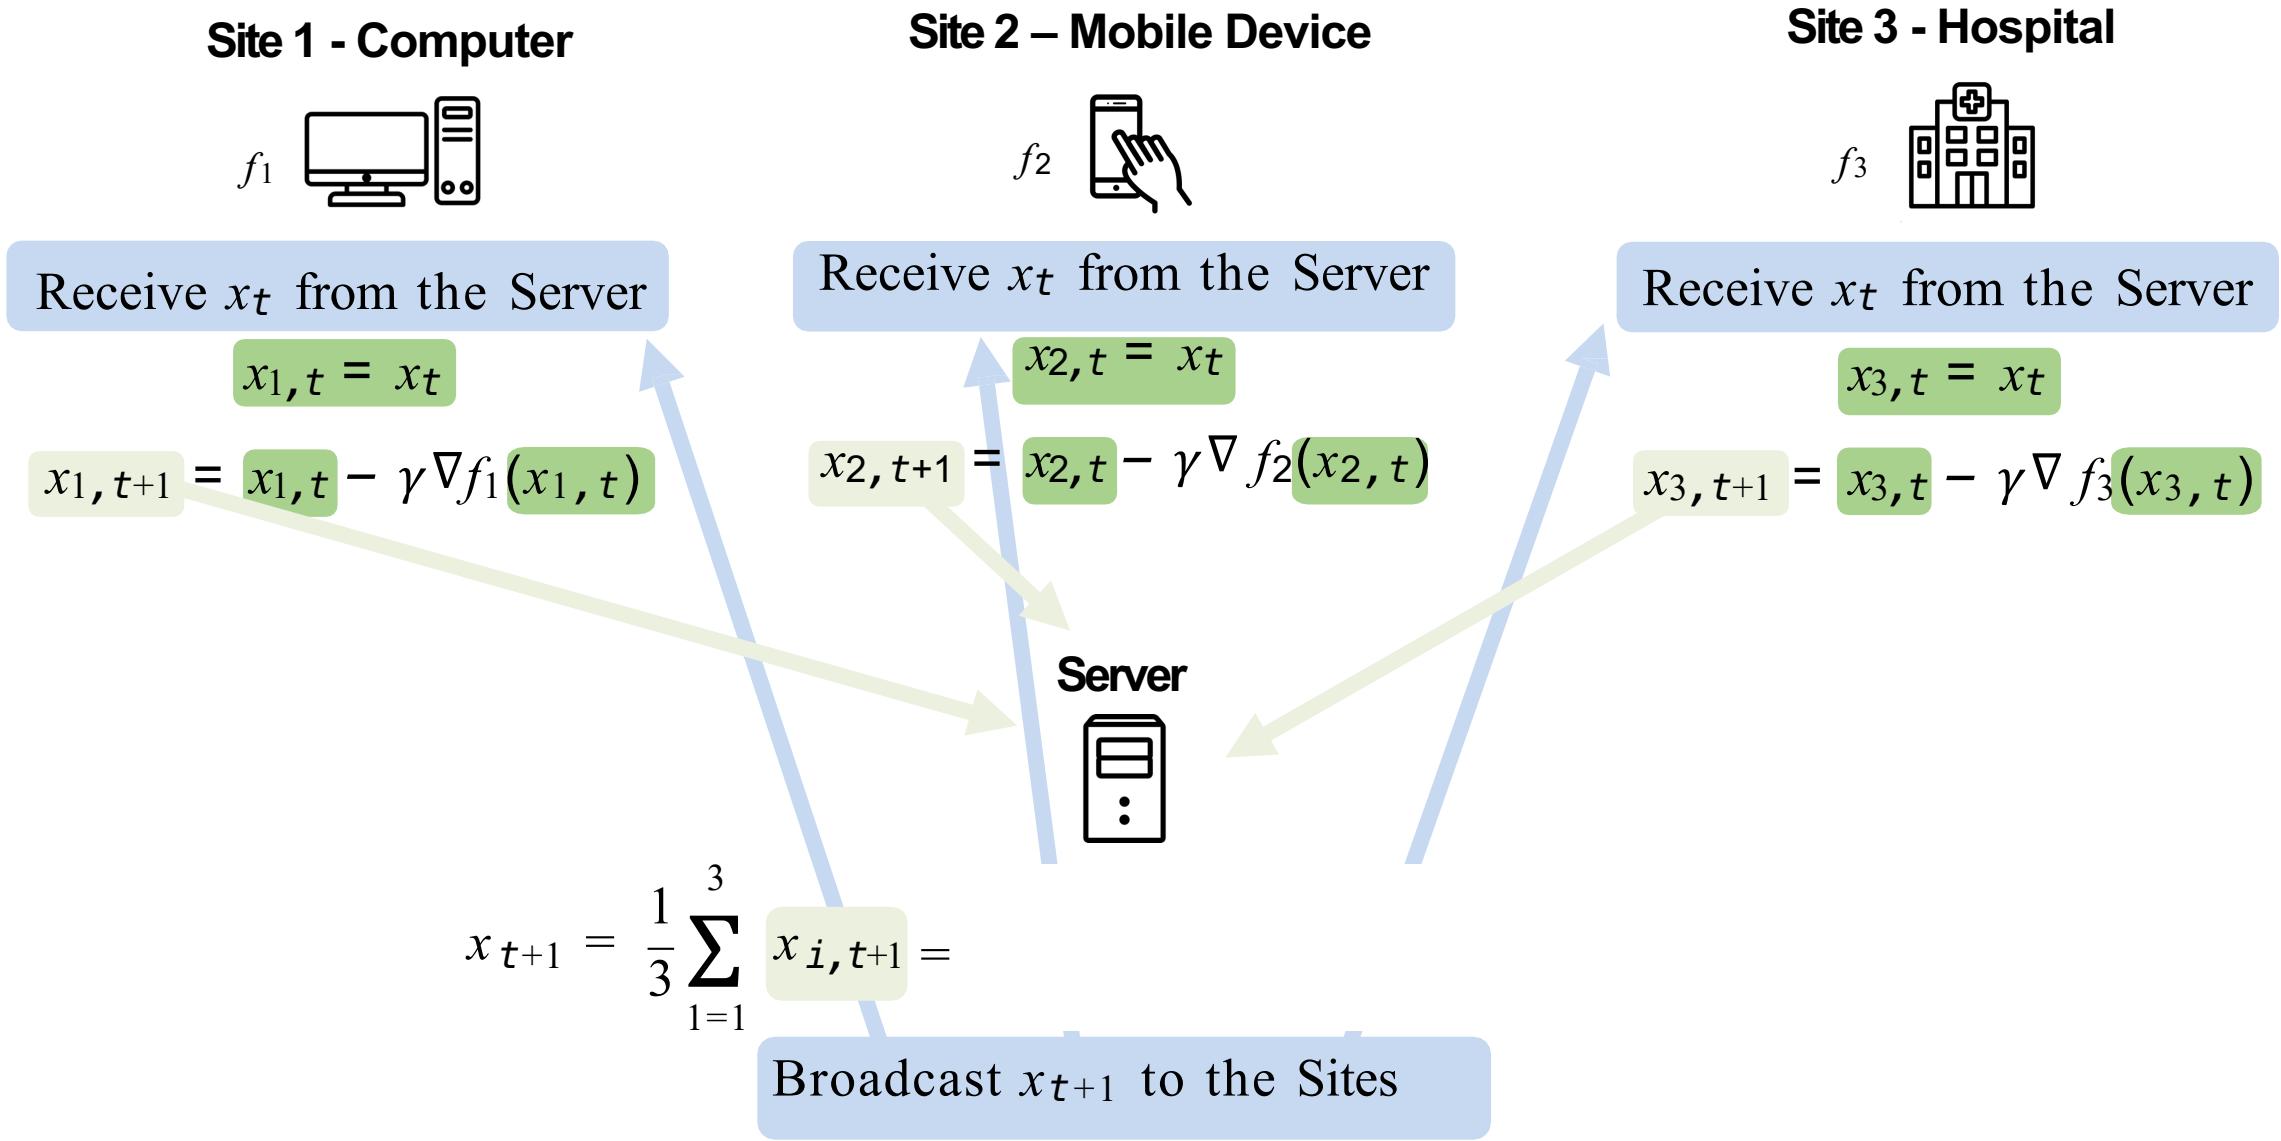

Distributed Local Gradient Descent

Site 1,2,3 – Computer, Mobile Device, Hospital

f_1 

Receive x_t from the Server

$$x_{1,t} = x_t$$

$$x_{1,t+1} = x_{1,t} - \gamma \nabla f_1(x_{1,t})$$

$$x_{1,t+K} = x_{1,t+K-1} - \gamma \nabla f_1(x_{1,t+K-1})$$

$$x_{t+K} = \frac{1}{3} \sum_{1-1}^{3} x_{i,t+K}$$

Broadcast x_{t+K} to the Sites

Server
Central Orchestrator

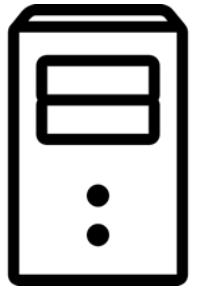

Part 3
Brief History of Local Training

From Gradient Descent to Local Gradient Descent

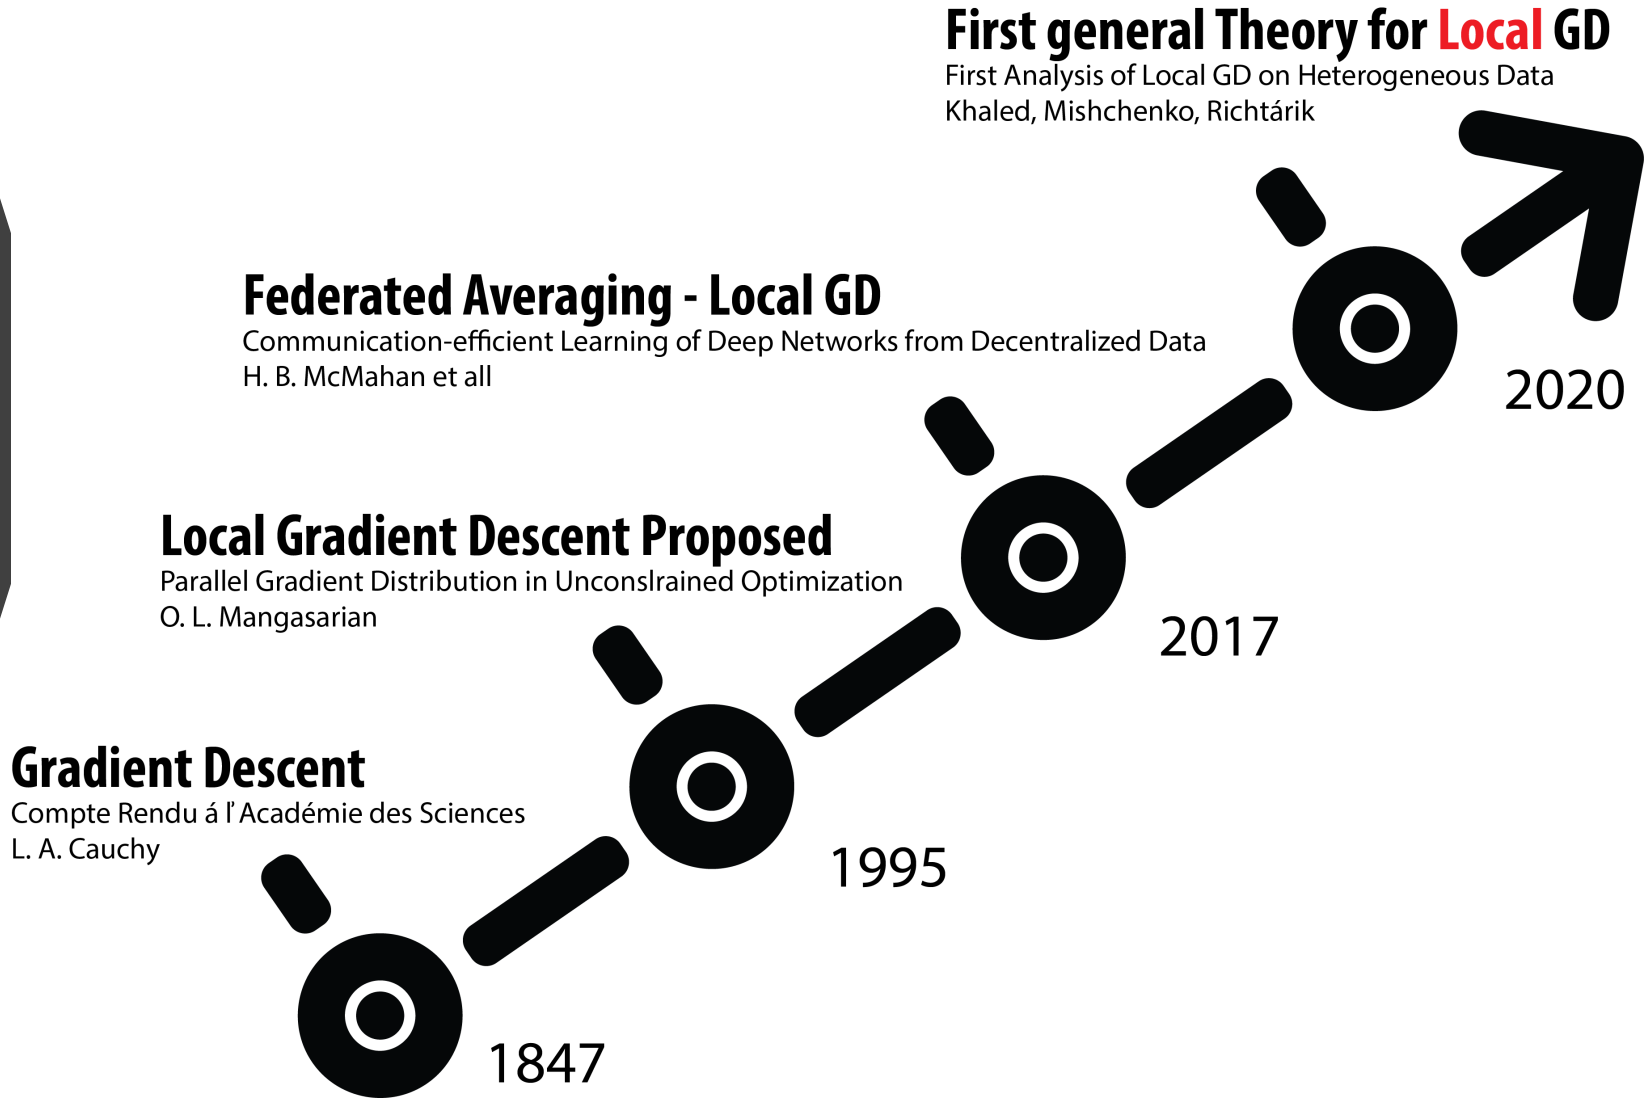

Part 4
What does Local Training do?

Local Training

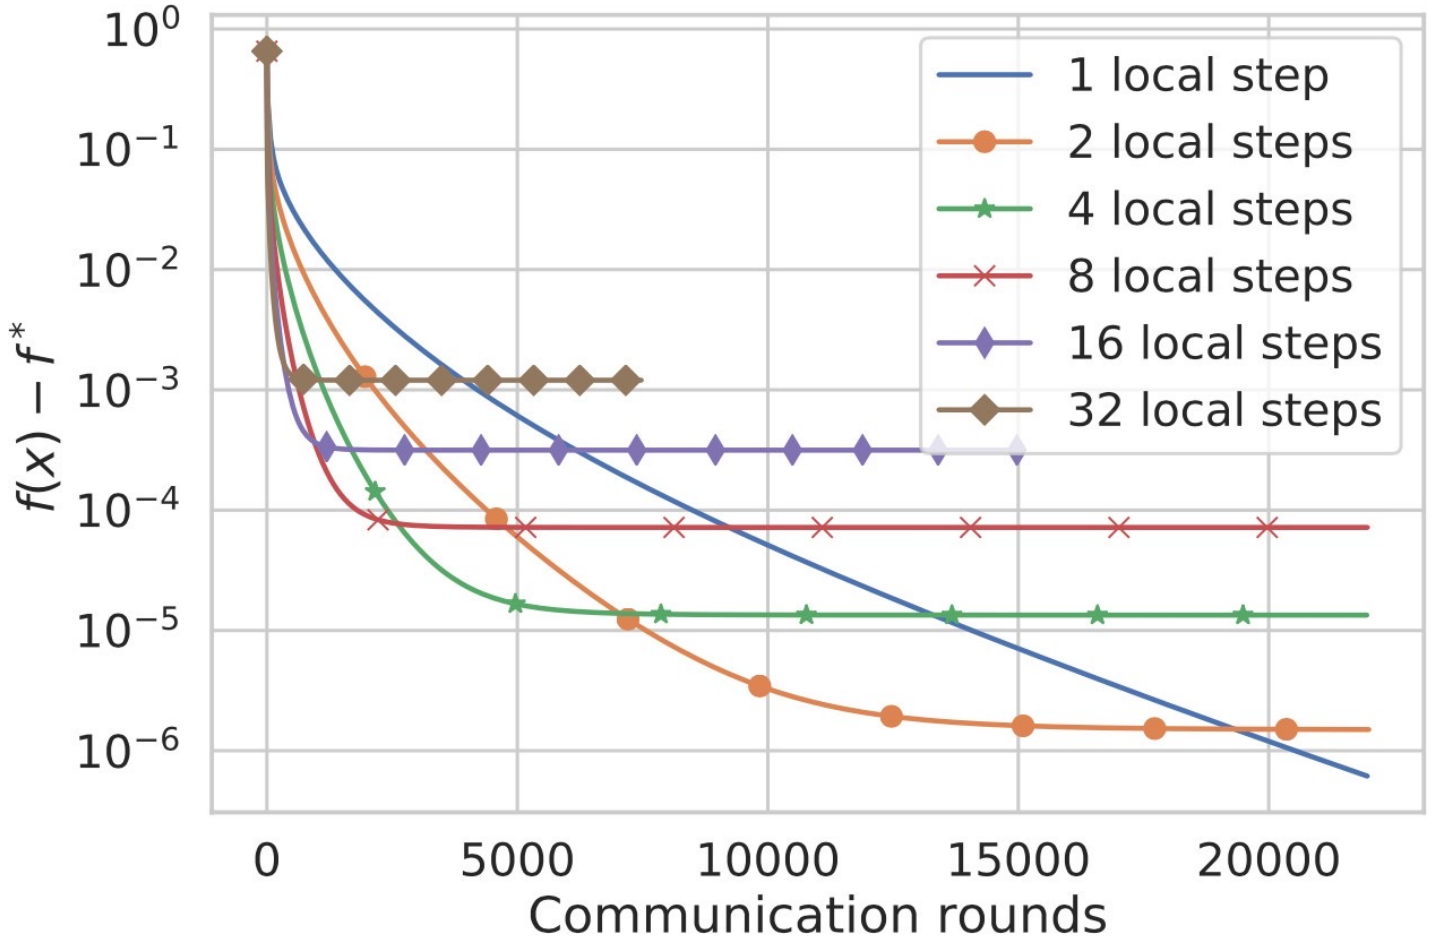

Thank You